

REMARKS

Claims 1-22 were pending. All were rejected. The applicants have cancelled claims 1-22 and added new claims 23-44. Therefore, claims 23-44 are presently pending. The applicants request further consideration and re-examination in view of the amendments above and remarks set forth below.

The applicants amended the specification to insert a new paragraph which describes Figure 4D. Support for this new paragraph can be found at least in Figure 4D and in original claim 6. Thus, no new matter has been entered.

Rejections under 35 U.S.C. § 112:

Claims 1-22 were rejected under 35 U.S.C § 112. The applicants have cancelled claims 1-22.

Rejections under 35 U.S.C. § 103:

Claims 1-22 are rejected as being unpatentable over U.S. Patent No. 6,421,809 issued to Wuytack (hereinafter “Wuytack”). The applicants have cancelled claims 1-22.

New Claims 23-44:

The applicants have added new claims 23-44.

New claim 23 recites “[a] method for computing assignments of data stores to storage device logical units for a data storage system design comprising: selecting a plurality of data stores for each of a plurality of data store clusters; clustering the selected data stores thereby forming the data store clusters; and assigning the data store clusters to the storage device logical units using a machine-implemented design algorithm by which constraint calculations are performed at the cluster level for each data store cluster.” Support for new claim 23 can be found at least in Figures 1, 2 and 3, and the descriptions thereof, and on page 3, line 2 to page 4, line 5 of the applicant’s specification.

The Wuytack reference is directed to memory-related design issues in the design of devices which include a memory organization and registers and for which operation of the devices can be described as a set of data access operations or instructions on data structures or variables being stored in said memories. Wuytack at col. 1, lines 13-20. Taken as a whole, Wuytack is clearly directed to design issues

related to volatile, short-term memory for processor-based computing systems. For example, Figure 2 of Wuytack shows the architecture of a digital device. The device includes a hierarchical distributed memory architecture, data paths, a global controller and address generators for the memories. Wuytack at col. 7, lines 52-55.

In contrast, the present invention is directed toward data storage system design. This is clear from new claim 23 and from the applicants' specification. For example, new claim 23 recites a method for computing assignments of data stores to storage device logical units for a data storage system design. A store represents a consumer of space on a storage device, such as a file, database, database table or file system. Applicants' specification at page 1, lines 13-15. A logical unit is a unit of storage having a particular size that is selected by taking into account the capacity of the storage device. Applicants' specification at page 1, lines 14-16. A type of storage device is a hard disk or hard disk array. See applicants' specification at page 7, line 6, (referring to RAID – a redundant array of independent or inexpensive disks) and at page 12, lines 9-10 (referring to array bandwidth).

“In order to rely on a reference as a basis for rejection of an applicant's invention, the reference must either be in field of applicant's endeavor or, if not, then be reasonably pertinent to the particular problem with which the inventor was concerned.” Manual of Patent Examining Procedure (MPEP), Section 2141.01(a) (Rev. 2, May 2004). In this same Section, the MPEP discusses *Wang Laboratories, Inc. v. Toshiba Corp.*, 993 F.2d 858, 26 USPQ2d 1767 (Fed. Cir. 1993), in which:

Patent claims were directed to single in-line memory modules (SIMMS) for installation on a printed circuit motherboard for use in personal computers. Reference to a SIMM for an industrial controller was not necessarily in the same field of endeavor as the claimed subject matter merely because it related to memories....

Manual of Patent Examining Procedure (MPEP), Section 2141.01(a) (Rev. 2, May 2004). Thus, a reference is not necessarily in the same field of endeavor as claimed subject matter merely because both relate to computer memories. Here, the applicants' claims are directed to storage system design, whereas the Wuytack reference relates to design issues related to volatile, short-term memory for processor-based computing systems. These are entirely different fields of endeavor; the Wuytack reference is simply not reasonably pertinent to the particular problem to

which the applicant was concerned - data storage system design. For at least this reason, claim 23 is allowable over Wuytack.

Further, even if the Wuytack reference could be applied, it does not teach or suggest all of the limitations of new claim 23. The portion of Wuytack which was cited by the examiner (col. 1, line 62 to col. 2, line 11) discusses a parallel compiler technique for SIMD (single instruction, multiple data) architectures, in which groups of data are assigned to different memories such that they can be accessed simultaneously. This technique is said to avoid allocation of multiport memories by explicitly splitting data arrays into smaller arrays that can be assigned to single port memories. Wuytack at col. 2, lines 5-9. This passage of Wuytack simply does not suggest or disclose the steps of selecting a plurality of data stores for each of a plurality of data store clusters; clustering the selected data stores thereby forming the data store clusters; and assigning the data store clusters to the storage device logical units using a machine-implemented design algorithm by which constraint calculations are performed at the cluster level for each data store cluster, as are recited by new claim 23. This is another reason why claim 23 is allowable over Wuytack.

New claim 24 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores according to an order in which the data stores are presented. Support for this feature can be found at least in Figure 5A and at page 11, lines 6-7 of the applicants' specification. New claim 24 is allowable at least because it is dependent from an allowable base claim 23.

New claim 25 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores in random order. Support for this feature can be found at least in Figure 5B and at page 11, lines 8-11 of the applicants' specification. New claim 25 is allowable at least because it is dependent from an allowable base claim 23.

New claim 26 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores according to an order in which the data stores are sorted based on characteristics of the data stores. Support for this feature can be found at least in Fig. 5C and at page 11, lines 12-15 of the applicants' specification. New claim 26 is allowable at least because it is dependent from an allowable base claim 23.

New claim 27 recites that the characteristics include size of the stores. New claim 28 recites that the characteristics include aggregate stream demands. Support

for these features can be found at least at page 11, lines 14-15 of the applicants' specification. New claims 27 and 28 are allowable at least because they are dependent from an allowable base claim 23.

New claim 29 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores according to an order in which the data stores are sorted based on similarity of characteristics. Support for this feature can be found at least in Figure 5D and at page 11, lines 16-22 of the applicant's specification. New claim 29 is allowable at least because it is dependent from an allowable base claim 23.

New claim 30 recites that the characteristics include reliability or storage class requirements. New claim 31 recites that the characteristics include layout requirements. New claim 30 recites that the characteristics include importance. Support for these features can be found at least at page 11, lines 17-19 of the applicants' specification. New claims 30, 31 and 32 are allowable at least because they are dependent from an allowable base claim 23.

New claim 33 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores according to an order in which the data stores are sorted based on desired properties for a particular data store cluster. Support for this feature can be found at least in Figure 5E and at page 11, line 23 to page 12, line 1 of the applicant's specification. New claim 33 is allowable at least because it is dependent from an allowable base claim 23.

New claim 34 recites that clustering the selected data stores for a particular one of the data store clusters comprises clustering a determined number of the data stores together. Support for this feature can be found at least in Figure 4A and at page 8, line 22 to page 9, line 5 of the applicant's specification. New claim 34 is allowable at least because it is dependent from an allowable base claim 23.

New claim 35 recites that clustering the selected data stores is based on capacity. Support for this feature can be found at least in Figure 4B and at page 9, lines 13-14 of the applicants' specification. New claim 35 is allowable at least because it is dependent from an allowable base claim 23.

New claim 36 recites that aggregate capacity requirements of the data store clusters prevent more than a particular number, N, of the data store clusters from being assigned per logical unit. Support for this feature can be found at least in Figure 4B and at page 9, line 13 to page 10, line 1 of the applicant's specification.

New claim 36 is allowable at least because it is dependent from an allowable base claim 23.

New claim 37 recites that clustering the selected data stores is based on bandwidth. Support for this feature can be found at least in Figure 4C and at page 10, lines 3-4 of the applicant's specification. New claim 37 is allowable at least because it is dependent from an allowable base claim 23.

New claim 38 recites that aggregate bandwidth requirements of the data store clusters prevent more than a particular number, N, of the data store clusters from being assigned per logical unit. Support for this feature can be found at least in Figure 4C and at page 10, lines 2-25 of the applicant's specification. New claim 38 is allowable at least because it is dependent from an allowable base claim 23.

New claim 39 recites that clustering the selected data stores is based on capacity and bandwidth wherein aggregate capacity and bandwidth requirements of the data store clusters prevent more than a particular number, N, of the data store clusters from being assigned per logical unit. Support for this feature can be found at least in Figure 4D of the applicant's specification and in original claim 6. New claim 39 is allowable at least because it is dependent from an allowable base claim 23.

New claim 40 recites that selecting a plurality of data stores for a particular one of the data store clusters comprises selecting the data stores in random order. Support for this feature can be found at least at page 12, lines 4-13 of the applicants' specification. New claim 40 is allowable at least because it is dependent from an allowable base claim 23.

New claim 41 recites de-clustering the data stores. Support for this feature can be found at least in Figure 3 and at page 7, line 25 to page 8, line 1 of the applicants' specification. New claim 41 is allowable at least because it is dependent from an allowable base claim 23.

New claim 42 recites determining whether all of the data stores have been assigned and, if all of the data stores have been assigned, outputting a resulting assignment. Support for this feature can be found at least at page 8, lines 1-3 of the applicants' specification. New claim 42 is allowable at least because it is dependent from an allowable base claim 23.

New claim 43 recites determining whether all of the data stores have been assigned and, if not all of the data stores have been assigned, reducing a cluster size and repeating said steps of selecting, clustering and assigning. Support for this

feature can be found at least at page 8, lines 3-21 of the applicants' specification. New claim 43 is allowable at least because it is dependent from an allowable base claim 23.

New claim 44 recites a computer-readable medium having stored thereon a computer program for performing a method of designing storage systems, the method comprising steps of: selecting a plurality of data stores for each of a plurality of data store clusters; clustering the selected data stores thereby forming the data store clusters; and assigning the data store clusters to the storage device logical units using a machine-implemented design algorithm by which constraint calculations are performed at the cluster level for each data store cluster. Support for new claim 44 can be found at least in Figures 1, 2 and 3, and the descriptions thereof, and on page 3, line 2 to page 4, line 5 of the applicant's specification and in original claim 22.

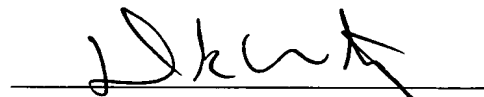
Thus, new claim 44 recites a method including steps of selecting, clustering and assigning, similarly to new claim 23. As discussed above, the Wuytack reference does not suggest or disclose these features. Thus, new claim 44 is allowable.

Conclusion:

In view of the above, the applicants submit that all of the pending claims are now allowable. Allowance at an early date would be greatly appreciated. Should any outstanding issues remain, the examiner is encouraged to contact the undersigned at (408) 293-9000 so that any such issues can be expeditiously resolved.

Respectfully Submitted,

Dated: Sept. 14, 2004

A handwritten signature in black ink, appearing to read 'Derek J. Westberg', written over a horizontal line.

Derek J. Westberg (Reg. No. 40,872)